ENTERPRISE INTELLIGENCE:
A NEW FRONTIER FOR INNOVATION

workday
Built for the future.®

I. Introduction

Big data analytics (BDA)—the use of data sets to gain actionable insights—and artificial intelligence (AI)—where computers simulate human behavior, including decision-making—are quickly changing our world. While coverage often focuses on futuristic applications of AI like robotic surgery and driverless taxis, AI and BDA in fact already power products and services we use every day, including a vast array of mobile apps, web services, and business solutions. They cut across virtually all sectors of the economy, transforming healthcare, transportation, logistics, entertainment, education, agriculture, and finance. In wearable devices, for example, these technologies help us live healthier lives by helping us track our steps, calories burned, and heart rate. They have revolutionized how we commute, underlying popular ride-hailing and ride-sharing apps. They cater to our tastes with services that offer us personalized music and video recommendations across numerous platforms. They identify aberrations in data in order to guard against bank fraud, keeping our money safe. And they keep track of how quickly students are learning course material, and how much water crops are getting.

While there is no shortage of coverage on the impact that BDA and AI will have on consumers, these technologies are likely to have a more immediate and profound impact on businesses and government. Companies of all sizes are beginning to use BDA and AI throughout the enterprise to streamline manufacturing, upgrade customer service, tighten supply chains, improve employee hiring and retention, and increase efficiency and productivity. Although uptake has been slower in the public sector, forward-looking governments are using BDA and AI to improve policing, mitigate traffic, facilitate access to public services, and generally to serve their constituents more effectively.

Given the pace of innovation, it is not surprising that policymakers are beginning to take note. To date, most policy discussions have centered on issues that directly affect consumers, such as product safety, liability, and privacy. Too often, however, these discussions use the term "artificial intelligence" loosely to cover many different technologies. They also fail to distinguish between business-to-consumer ("B2C") and business-to-business ("B2B") applications and the distinct social and policy issues that each raises. In contrast to the discourse that surrounded past technology transformations (such as the widespread adoption of personal computing and the internet), today's policy discussions also often focus heavily on possible risks, and less on the vast promise that these technologies hold for people and societies.

As these discussions move forward, a key challenge will be to ensure that well-meaning efforts to protect consumers from harm do not inadvertently frustrate innovation or hinder beneficial uses of AI and BDA. As a first step, stakeholders need to work toward a common vocabulary based on a clear understanding of the technologies involved. They also need to distinguish between applications that directly affect consumers and those that are used exclusively in the enterprise context, given the different policy concerns that each of these scenarios may raise.

Perhaps most importantly, we must bear in mind that these technologies are not an autonomous force beyond our control. All of us—developers, users, consumers, and policymakers—have the power to determine how these technologies evolve, how they are applied, and ultimately how they impact people and societies. Through open discussion and engagement, we can, and indeed must, promote the thoughtful and responsible adoption of AI. As one of the leading providers of data analytics technologies for enterprises, Workday takes this responsibility seriously, and offers this paper in an effort to advance an understanding of these issues.

Part II of this paper begins by describing several technologies that often are lumped under the term "artificial intelligence," then provides case studies of recent applications of AI and BDA across a range of sectors and use scenarios. Part III surveys recent efforts by governments to grapple with the policy issues that these technologies raise and also examines recent industry and other private-sector initiatives in this area. Part IV draws on these efforts as a basis for offering high-level principles for how the public and private sectors might work together to approach the policy issues raised by BDA and AI. We also offer a proposal for an industry code of conduct.

The conclusion calls for a dialogue among policymakers, regulators, civil society, academics, and industry to continue developing and refining approaches for how best to leverage the promise of BDA and AI while also promoting the ethical and responsible use of these technologies across society.

II. Technology Overview and Case Studies

A. Technology Drivers and Key Terms

Despite the current media hype, the idea of artificial intelligence is not new. AI research has existed for decades, and discrete applications of AI have appeared in business and government for years. But in the past few years, a confluence of technology and industry developments has shifted BDA and AI innovation into high gear. These advances include: (1) the wide uptake of devices and services that generate massive amounts of data, such as smartphones, sensors, connected devices, and interactive online services; (2) plummeting costs for data storage; (3) the emergence of cloud computing, which has made massive computer power available to anyone with a broadband connection and a credit card; and (4) rapid advances in data analytics software.

In this new data-driven economy, "big data" and "AI" have become buzzwords. But while these terms are widely used, there is a lack of clarity on what they mean, even among experts. AI in particular often means different things to different people, sweeping in machine learning, robotics, data analytics, augmented intelligence, and a range of other technologies. Understanding the differences between them, however, is critical to informed policymaking, since each offers a different mix of potential benefits and risks.

In hopes of promoting greater clarity in this area, the points below offer proposed working definitions of certain key terms in this area:

- **BIG DATA**. The big data phenomenon is a foundational element of current advances in BDA, AI, and computational analytics, more broadly. The term typically refers to large data sets, either structured or unstructured, that are analyzed to reveal patterns, trends, and associations that may be invisible from the analysis of smaller or more uniform data sets. The technology trends identified above—including the ability to collect huge amounts of data from a variety of sources and the increasing affordability of data storage—have contributed to the rise of big data.[1]

[1] See *Big Data: A Tool for Inclusion or Exclusion? Understanding the Issues*, Fed. Trade Comm'n. (Jan. 2016), https://www.ftc.gov/system/files/documents/reports/big-data-tool-inclusion-or-exclusion-understanding-issues/160106bigdata-rpt.pdf, at 2 [hereinafter "FTC Big Data Report"].

- **BIG DATA ANALYTICS**. Big data analytics refers to the analysis of big data, using sophisticated algorithms, to inform human decision-making. The analysis of data pools can provide instantaneous and useful insights, uncover hidden correlations, and allow individuals and businesses to make better decisions.[2] In contrast to AI (described below), BDA typically does not involve automated decision-making. Rather, by using data to reveal connections and insights that were previously invisible, BDA provides information that empowers people to make better decisions. The end result of BDA is the provision of real-time, serviceable information that is actionable by people.

Workday Benchmarking

Workday Benchmarking ("Workday Benchmarking") is the first product offering delivered under the Workday Data-as-a-Service platform and is a great example of BDA in action. Workday customers can utilize Workday Benchmarking to leverage the collective power of the extensive Workday community—including over 26 million workers across more than 1,800 global organizations—by contributing selected de-identified data to a shared, aggregated, and secure data set for insights that will help drive more informed decision-making. Each Workday customer that participates in Workday Benchmarking gains access to reports and analysis that provide key metrics to better understand each customer's relative performance in comparison to their peers.

Workday Benchmarking currently gives participating customers access to a variety of benchmarks such as workforce composition (age, diversity, tenure), turnover and career retention (talent, turnover, and others), leadership and manager effectiveness (span of control and leadership), and Workday usage (system utilization, business processes). In the near future, Workday will roll out financial management benchmarks that will include reports on core revenue growth, return on invested capital, and others. Each benchmark is surfaced in Workday applications and dashboards so that business leaders are empowered to make informed decisions to improve their organization's competitive position without ever having to leave their system of record.

Traditional benchmarking services provide outdated data delivered out of context, which limits the analytical power of such services. Workday Benchmarking leverages the power of the Workday technology platform and participating customer base to ensure that all reports and analysis delivered to customers are reliable and accurate. The entire Workday Benchmarking data set is purged and rebuilt monthly so that participating customers can gain insights and make decisions based on then-current market conditions. With access to current and reliable data that spans a cross-section of similarly situated peers, each Workday customer has the tools to achieve optimal performance in their respective markets.

[2] Adapted definition from *Merriam-Webster*, which defines "analytics." See *Merriam-Webster*, Analytics, available at https://www.merriam-webster.com/dictionary/analytics (last visited Oct. 10, 2017).

- **ARTIFICIAL INTELLIGENCE.** AI is commonly understood as technology that enables computers to simulate intelligent human behavior. To quote *Oxford Dictionaries,* it refers to "the theory and development of computer systems able to perform tasks normally requiring human intelligence…"[3] Like BDA, many AI technologies today rely on complex algorithms to analyze massive data sets, from which they identify patterns and then draw inferences and conclusions. Unlike BDA, however, most applications of AI do entail the training of computers to complete certain tasks in an automated way. For instance, AI applications may decide when to execute turns in self-driving cars, when to adjust thermostats, or when to water crops. In this way, AI technologies not only inform human decision-making, they can to some extent make decisions and undertake actions in place of humans. BDA, by contrast, yields insights that help decision-makers analyze, answer, and act—but does not act as the decision-maker itself.

  Using these definitions, Workday Benchmarking would qualify as an example of BDA rather than AI because the service does not make actual decisions or undertake actions in place of humans. Instead, Workday Benchmarking reveals trends or patterns that customers can then use to inform their decision-making.

  It is worth emphasizing, however, that even AI applications remain fully within the control of human beings—despite what many current science fiction movies would have us believe. For that reason, experts sometimes use the term "augmented data discovery" in place of "artificial intelligence," since that underscores the role human intelligence and control play even within AI applications. While a sophisticated AI program is certainly capable of making decisions, those decisions are limited to those determined by humans, and are only as good as the data used to train the program.[4]

- **ALGORITHMS.** Algorithms form the computational processes that power BDA, AI, and many related technologies. In their simplest form, algorithms constitute a set of step-by-step instructions that computers follow to perform a particular task.[5] Algorithms are used for calculation, data processing, and automated reasoning, and enable technological solutions that enhance perception, learning, and decision-making aimed at improving the ability of people

[3] Oxford University Press, *Oxford Living Dictionaries* (2017), *available at* https://en.oxforddictionaries.com/definition/artificial_intelligence.

[4] Margaret Rouse, *What Is Augmented Intelligence*, http://whatis.techtarget.com/definition/augmented-intelligence (last visited Oct. 14, 2017).

[5] See *Artificial Intelligence: Opportunities and Implications for the Future of Decision-Making*, UK Gov't Office for Sci. (2015), https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/566075/gs-16-19-artificial-intelligence-ai- report.pdf, at 5 [hereinafter "UK AI Report"].

to solve complex problems. Effectively all of the technologies listed below use algorithms to process information.

- **MACHINE LEARNING.** Machine learning refers to AI algorithms that change in response to their own output—perhaps better described as "computer programs that automatically improve with experience."[6] Machine learning systems often identify difficult-to-spot relationships in data that may otherwise have been missed.[7] Most machine learning approaches are not restricted to producing a single prediction from given inputs; instead, they often produce probabilistic outputs, offering a range of likely predictions with associated estimates of uncertainty.[8]
- **DEEP LEARNING/NEURAL NETWORKS**. Deep learning is a subset of machine learning. Also known as multilayered neural networks, it is a process that depends on using layers of non-linear algorithmic processes to find patterns or classify data.[9] Although this general approach includes a variety of different techniques, the key feature is that they each use a layered or staged design, in which the outputs from the previous layer are used as inputs for the next.[10]

Employee Development Powered by Machine Learning

More and more, companies are focusing on the development of individual employees. Workday supports this through our machine learning capabilities in several ways, including the following:

Workday Learning applies machine learning to tailor workplace learning to the individual; i.e., the system knows the employee and makes course and lesson recommendations based on things such as the worker's position, tenure at the company, interaction with content, and so on. In a dynamic business environment where automation and machine learning are evolving jobs, companies need to readjust learning strategy and programming to ensure employees learn new skills, continue to grow in their roles, and prepare for new challenges.

The career paths/opportunity graph is a feature of the core Workday Human Capital Management system that enables everyone to take charge of career planning and proactively create career goals. It gives employees a personalized view of career paths within their companies by revealing the job moves of others who have held the same role. By analyzing the skills of those individuals, it helps employees focus on building the right skill set for targeted roles. As a result, employees are empowered to develop their careers within—instead of outside of—the company.

[6] UK *AI Report*, supra note 1, at 6 (internal citations omitted).

[7] See *id.*

[8] See *id.*

[9] See *id.* at 7.

[10] See *id.*

- **ROBOTICS**. Robotics generally refers to the implementation of AI into devices that perform certain tasks at least somewhat autonomously (although some forms of robotics, of course, do not include any element of AI). Self-driving cars and grocery delivery vehicles[11] are common examples of robotics in the consumer space, but industrial enterprises have also begun leveraging robotics to implement more streamlined processes in complex manufacturing spaces, such as the automotive and aerospace industries.[12]

As noted, while all of these technologies typically analyze large and diverse data sets to detect new patterns and relationships, they differ on whether the outputs of this analysis are used to inform human decision-making or instead to make decisions or act on behalf of humans. BDA stops with the former; AI and robotics typically engage in the latter. Through thoughtful engineering and policymaking, all of these technologies can continue to grow in an ethical and responsible way. But these distinctions must remain front of mind when considering the balance of benefits and risks that each poses, and the corresponding need for regulation.

B. Benefits of Big Data and AI for Enterprises

In addition to a clear understanding of the technologies and terms involved, effective policymaking should recognize, and seek to advance, the potential benefits of these technologies. Although we are in the early stages of BDA and AI innovation and adoption, it is clear that both have the potential to generate enormous opportunities and advantages for enterprises.

BDA and AI are enabling companies across virtually all sectors of the economy to become more efficient, productive, and innovative. Enterprises are already using solutions powered by both technologies to help streamline logistics, improve inventory management, better manage customer relations, and hire, retain, and serve employees more effectively. For instance, by enabling enterprises to continuously track and analyze customer data, BDA can help businesses acquire insights they can use to improve the customer experience and customer satisfaction. As Workday Benchmarking illustrates, BDA can also provide broader insights into how a

[11] Kat Londsdorf, *Hungry? Call Your Neighborhood Delivery Robot*, NPR (Mar. 23, 2017), available at http://www.npr.org/sectionsalltechconsidered/2017/03/23/520848983/hungry-call-your-neighborhood-delivery-robot.

[12] Daniel Faggella, *Global Competition Rises for AI Industrial Robotics*, TechEmergence (Sept. 14, 2017), available at https://www.techemergence.com global-competition-rises-ai-industrial-robotics/.

company's performance, growth, and retention rates compare with other similar organizations.

Workday Prism Analytics

Workday Prism Analytics is another example of BDA in action. Workday Prism Analytics is a self-service solution that allows participating Workday customers the ability to bring together data from a variety of different sources (including data from the Workday technology platform) and apply leading-edge analytics tools to gain insights into financial and people metrics in order to optimize performance throughout an organization.

To stay competitive in any industry, business leaders need to make informed decisions based on data presented from a variety of sources. Historically, organizations had to pull data from many different sources and systems of record and then shift the data into separate systems to analyze, visualize, and report. This process was a disjointed patchwork of systems, security models, and experiences, which often led to decision-makers receiving analysis that was stale, static, unreliable, and received outside of the finance and HR systems that organizations relied on to run their business.

Workday Prism Analytics has redefined the data analysis process by allowing customers to upload data from many different sources to analyze finance and HR business questions right in the context of their system of record. With Workday Prism Analytics, customers have access to a wide range of use cases, such as organizational health, financial forecasting, worker productivity, and revenue. For example, business leaders in the technology space can view regional attrition data from the Workday technology platform together with third-party survey data on employee engagement globally in order to better manage employee turnover. By analyzing current and reliable data from a variety of sources in a single platform, Workday Prism Analytics customers can make informed decisions to optimize business performance across their organization.

Enterprises are similarly using AI to improve business performance. Companies can use AI to accelerate their production capabilities through more reliable demand forecasting and increased flexibility in operations and supply chains. In doing so, they can create smarter, faster, cheaper, and more environmentally friendly production processes that increase worker productivity, refine product quality, lower costs, and improve worker health and safety. Enterprises can also use supply AI to improve supply-chain management via adaptive scheduling and routing. AI helps supply chains become more secure from disruption by making automatic adjustments for the anticipated effects of weather, traffic, and unforeseen events. For instance, AI-powered distributed sensor systems can analyze environmental conditions and

detect when the probability of major infrastructure disruptions increases significantly, helping enterprises adapt operations as needed to respond to disruptions even before they occur.[13] Relatedly, BDA and AI offer tremendous promise in the context of employment. BDA can be used to uncover and possibly reduce employment discrimination by better enabling companies to objectively consider experiences and skill sets that have a proven correlation with success. By examining the skills that have made previous employees successful, BDA can help human resources officers more effectively "pattern match" in order to recognize the characteristics the next generation of hires should have.[14] In addition, as explained further in Part IV below, the growth of big data and AI is creating significant new opportunities for employment, particularly in research and development fields.

C. Case Studies

When people think about BDA and AI, they often think first of consumer-facing applications, such as virtual personal assistants or autonomous vehicles. But BDA and AI are already transforming enterprise productivity and the public sector, in particular education, in ways that include the following:

- **EMPLOYEE RETENTION**. The United States has recently seen some of the highest job opening and turnover rates in years—about 5.1 million openings and 2.7 million voluntary resignations, according to the latest JOLTS report.[15] With the median cost of turnover at about 20 percent of an employee's annual salary, losing top performers affects more than just quality of work and customer satisfaction: it can have a significant impact on a company's bottom line.[16] As a result, companies across industries around the world are looking for new approaches to people management and company culture to increase employee retention and satisfaction.

[13] See U.S. Executive Office of the President, *The National Artificial Intelligence Research and Development Strategic Plan*, National Science and Technology Council (Oct. 2016), https://nitrd.gov/PUBS/national_ai_rd_strategic_plan.pdf, at 19 [hereinafter "NSTC AI Strategic Plan"] (citing Bela Genge, Christos Siaterlis, and Georgios Karopoulos, Data fusion-based anomaly detection in networked critical infrastructures, 43rd Annual IEEE/IFIP Conference on Dependable Systems and Networks Workshop [DSN-W], 2013).

[14] See *U.S. Executive Office of the President, Big Data: A Report on Algorithmic Systems, Opportunity, and Civil Rights* (May 2016), https://obamawhitehouse.archives.gov/sites/default/files/microsites/ostp/2016_0504_data_discrimination.pdf at 14 [hereinafter "White House Big Data Report"].

[15] See *4 Questions on Tackling Retention Problems with Workday Talent Insights*, Workday (Apr. 9, 2015), available at http://blogs.workday.com/4-questions-on-tackling-retention-problems-with-workday-talent-insights/.

[16] See *id.*

Workday Retention Risk Analysis

Workday also provides BDA in the retention risk analysis feature of the core Workday Human Capital Management system. Retention risk analysis gives Workday customers the ability to quickly analyze their current roster of employees to identify and understand retention risk for any department, including the number of top performers at a high risk of leaving the organization within the next year.

Keeping top talent is a major priority for virtually every organization. With the median cost of employee turnover at about 20 percent of an employee's annual salary, employers that experience high employee turnover are not only losing out on continuity and quality of work—it is also dramatically affecting their bottom line. Historically, organizations have attempted to analyze employee turnover issues by some combination of hiring consultants, technology, and developing their own solutions. This disjointed approach is very expensive and time-consuming, and in many cases has not led to the type of insights that organizations need to address the issue.

Retention risk analysis can help by providing analysis and reports that highlight the likelihood that certain high-performing employees could leave an organization within a user-defined time frame. This is achieved by identifying risk factors for particular employees or even specific departments. For example, retention risk analysis looks at a number of variables, such as time in a current job function, number of locations where an employee works, number of job functions held, and even time between promotions, in order to provide analysis that predicts how many employees will leave their current jobs within a certain time frame and what the cost would be to the specific organization. Retention risk analysis can also recommend ways in which an organization can retain top talent by suggesting career moves for high-performing employees that have been identified as retention risks. These reports and analysis are invaluable considering that most medium- and large-size organizations could save tens of millions of dollars simply by retaining their best employees.

- **JOB MATCHING**. BDA and AI are helping place workers in positions that best align with their skills and strengths. For prospective employees, AI is being used to power tools that gather information about a particular candidate and the sort of position they seek, then using that information to develop a tailored list of recommended job listings. Some of these tools, for instance, can scan job postings from a network of *Fortune* 1000 companies and alert a user if they are a potential top-10 candidate for an opportunity.[17] If the user is interested, the program can forward an application to the prospective employer within 3 minutes.[18] These tools can help protect the privacy of job seekers by ensuring that the only entities that see the user are those companies to which a user submits an application.[19]

[17] See *Jobs at Stella*, https://angel.co/stella-3/jobs (last visited Oct. 19, 2017).

[18] See *id.*

[19] See Kayla Matthews, *5 chatbots that will help you find a job,* VentureBeat (June 22, 2017), available at https://venturebeat.com/2017/06/22/5-chatbots-that-will-help-you-find-a-job/.

AI tools can even help prospective employees craft more effective resumes.[20] They can assist in determining what traits companies are looking for and what keywords to include in the resume, assess the resume's strengths and weaknesses, and edit it accordingly.[21] These sorts of tools have, for instance, helped veterans find civilian jobs that correlate to skills they obtained in the military.[22]

For employers, AI tools can help drastically reduce the time needed to find a new hire. Employers have reported that these tools can help narrow a pool of 4,000 candidates to the top 2 to 3 percent within a few days,[23] and some tools are able to reduce search time by 80 percent—or from around two months to two weeks.[24] Furthermore, BDA tools can assist employers in identifying applicant traits that more accurately correlate to workplace success. For instance, BDA has revealed that, contrary to conventional wisdom suggesting that successful salespeople must have an outgoing and naturally friendly personality, salespeople are in fact more likely to succeed when they have emotional courage and persistence.[25] AI tools can also assist employers in posting job descriptions aimed at improving diversity and that can hide an applicant's name, gender, and personal identifiers to reduce opportunities for decisions based on unconscious biases.[26]

- **ENHANCING THE WORKER EXPERIENCE**. BDA and AI tools can also help improve an employee's professional experience once they are situated in a position. For instance, AI tools can provide employees easier and quicker access to common HR questions. One AI tool, for example, enables employees to ask a chatbot question, such as "Are we off on President's Day?" or "What are my dental benefits?" and simultaneously enables the chatbot to reach out to employees with information that

[20] See *id.*

[21] See *id.*

[22] See *id.*

[23] See Jennifer Alsever, *How AI Is Changing Your Job Hunt*, *Fortune* (May 19, 2017), available at http://fortune.com/2017/05/19/ai-changing-jobs-hiring-recruiting/.

[24] See *Matthews, supra* note 33.

[25] Rudy Karsan, *With Big Data, Companies Can Predict Your Success Before Your First Day On The Job*, FastCompany (July 26, 2013), available at https://www.fastcompany.com/3014837/with-big-data-companies-can-predict-your-success-before-yo.

[26] See Matthews, *supra* note 33.

may be of interest, such as extracurricular events the company is offering on a particular day.[27] This sort of automated information distribution can create more time for HR professionals to focus on broader and more complex questions about how to promote employees' professional development.

BDA can also assist employers in more proactively addressing employee concerns. For instance, BDA tools can track the frequency with which multiple employees are raising the same concern and thus help identify when a common, company-wide problem exists so that employers can offer solutions in a timely manner.[28]

- **WORKER VALUE**. BDA and AI tools can also enhance enterprise performance by enabling employees to focus on tasks to which they add the greatest value. BDA tools, for instance, can help employers more effectively understand drivers of employee productivity and reveal areas that could benefit from improvement. For example, big data can help identify causes of fatigue, such as sitting too long, and help address them. It also can identify how employees interact with each other and customers, helping to find natural teams.[29] In addition, studies have shown that increasing the usability of data within an organization—i.e., presenting data more concisely and consistently across platforms such as corporate laptops and mobile devices, and allowing it to be more easily manipulated—can yield billions in revenue increases. In one example, a University of Texas at Austin study found that the median *Fortune* 1000 business could see average revenue growth of $2.01 billion annually if it increased the usability of its data by just 10 percent.[30]

  Solutions powered by AI likewise offer tremendous opportunities to boost worker productivity. AI applications help automate lower-level activities and thereby free employees to focus on higher-value work or discover innovative ways to provide value.[31] In the cybersecurity context, for example, IBM's AI-based Watson for Cyber Security can analyze 15,000 security documents per day—a rate essentially impossible for any individual

[27] See Jeanne Meister, *The Future Of Work: The Intersection of Artificial Intelligence and Human Resources*, *Forbes* (Mar. 1, 2017), available at https://www.forbes.com/sites/jeannemeister/2017/03/01/the-future-of-work-the-intersection-of-artificial- intelligence-and-human-resources/#2ca6d1c96ad2.

[28] See *id.*

[29] *3 Ways Data Improves Productivity in the Workplace*, Innovation Enterprise (Jan. 20, 2017), available at https://channels.theinnovationenterprise.com/articles/3-ways-data-improves-productivity-in-the-workplace

[30] Anitesh Barua, Deepa Mani, and Rajiv Mukherjee, *Measuring the Business Impacts of Effective Data* (2015), available at http://www.datascienceassn.org/sites/default/files/Measuring%20Business%20Impacts%20of%20Effective%20Data%20I.pdf, at 3.

[31] *Employees: An Endangered Species?*, KPMG (Feb. 2016), available at http://www.kpmg-institutes.com/institutes/sharedservices- outsourcing-institute/articles/2015/09/employees-endangered-species.html, at 9.

to achieve. Watson thus enables analysts to build off of its work to more quickly identify incidents that require human attention.[32]

- **ENTERPRISE EFFICIENCY.** BDA and AI tools can also help generate efficiencies in enterprise management. GE, for instance, has used AI and machine learning to streamline its mechanical processes—including machine learning approaches that predict required maintenance for its large industrial machines. GE deploys AI technologies that use information from GE's machines in the field to generate and modify predictive models over time.[33] In addition, GE has also used AI and machine learning to analyze business data and to identify and normalize differential pricing in its supplier data across business verticals—leading to a savings of $80 million.[34]

Technologies powered by AI and BDA are also set to transform governments and the delivery of public information and services, particularly in education. A small sampling of examples of how governments are beginning to use these technologies in schools help illustrate their tremendous potential:

- **IMPROVING EDUCATIONAL OUTCOMES.** BDA and AI tools are transforming how schools are run and how educators teach students. These tools already analyze data produced by learning software and digital games, for instance, to offer insights about the pace and extent of student progress. They also power software that is adaptable to student needs, identify areas where courses need to improve, and break down textbook content into digestible "smart" study guides.[35] More broadly, BDA and AI solutions can help school districts gain insights on performance and attendance patterns.[36]
- **EMPOWERING BETTER CHOICES**. The U.S. government has used BDA tools to offer students and families reliable information about college performance and thereby position students to better evaluate college choices.[37] In particular, the U.S. Department of Education uses BDA to

[32] See *IBM Delivers Watson for Cyber Security to Power Cognitive Security Operations Centers*, IBM (Feb. 13, 2017), available at https://www-03.ibm.com/press/us/en/pressrelease/51577.wss; Jason Corbin, *Bringing the Power of Watson and Cognitive Computing to the Security Operations Center*, SecurityIntelligence (Feb. 13, 2017), available at https://securityintelligence.com/bringing-the-power-of-watson-and-cognitive-into-the-security-operationscenter/?cm_mc_uid=70595459933115020631816&cm_mc_sid_50200000=1503364089&cm_mc_sid_52640000=1503365578.

[33] Jayanth Kolla, *How Artificial Intelligence benefits companies and ups their game*, livemint (July 24, 2017), available at http://www.livemint.com/Technology/gLdRAxSjWVWCSJRl6E7bbL/How-firms-are-using-artificial-intelligence-to-up-theirgame.html.

[34] See *id.*

[35] See Daniel Faggella, *Examples of Artificial Intelligence in Education*, TechEmergence (Sept. 1, 2017), available at https://www.techemergence.com/examples-of-artificial-intelligence-in-education/.

[36] Benjamin Herold, *Are schools ready for the power and problems of big data?*, Education Week (Jan. 11, 2016), available at http://www.edweek.org/ew/articles/2016/01/13/the-future-of-big-data-and-analytics.html.

[37] See *White House Big Data Report*, *supra* note 14, at 17.

create a "College Scorecard," which furnishes students and their families with a more accurate picture of college cost and value and empowers them to make more informed decisions and better understand the opportunities and trade-offs of their educational choices.[38]

- **REDUCING DISPARITIES.** BDA can help students overcome disparities in learning outcomes and provide extra assistance for those more likely to drop out or fail. For instance, Georgia State University used BDA to develop a Graduation and Progression Success ("GPS") advising program designed to keep the school's more than 32,000 students on track for graduation.[39] The program tracks 800 different risk factors for each student on a daily basis, so that when a problem is detected, the university can deploy proactive advising and timely interventions to provide the support that student needs.[40] The results speak for themselves: over the past three years, Georgia State's graduation rate has increased by 6 percentage points, from 48 percent to 54 percent—with the biggest gains among those students in at-risk populations.[41]

As these examples demonstrate, technologies powered by AI and BDA are creating benefits across-the-board. Quite literally, they are transforming our world and how we live our lives.

III. Policy Issues and Potential Responses

In part because of their broad and potentially profound impact, BDA, AI, and related technologies are forcing thought leaders in both the public and private sector to grapple with the question of how to maximize the benefits of these technologies while minimizing their potential risks. This question, in turn, has prompted debate on several more-focused policy issues, such as:

- How can developers and users of these technologies prevent them from creating or perpetuating unintended biases?
- How do we ensure that AI technologies in particular treat people fairly and are applied in a transparent and accountable manner?
- How do we ensure that workers whose jobs are displaced by these technologies can acquire the skills needed for new jobs?
- What are the implications of these technologies for education, and how do we ensure that students acquire the skills that will be most in demand as a result of the AI revolution?

[38] See *id.*

[39] See *id.*

[40] See *id.*

[41] See *id.*

These are difficult questions, and answering them will often require making nuanced distinctions based on the technologies involved, the use scenarios at issue, and other factors. For instance, because BDA tools (as explained above) are primarily designed to enable *humans* to make better-informed decisions, and AI applications to some extent supplant human decision-making, the appropriate policy response might be different for each. Likewise, certain concerns—for instance, in relation to consumer safety and privacy—are likely to be less pressing for inventory management or other enterprise solutions than for devices or services that interact directly with consumers.

All of this suggests that a "one-size-fits-all" approach to policymaking in this area is unlikely to achieve the right balance. Instead, policy prescriptions—whether voluntary or mandatory—need to take account of material differences in technologies, applications, and use scenarios. Also, given that we are in the early stages of innovation in this area and much remains unknown as to how these technologies will develop, it is worth considering the valuable role that the private sector can play in creating flexible, context-specific best practices and "rules of the road" with regard to both the development and deployment of these technologies.

If we look across the policy landscape today, there are grounds for optimism. Most governments, recognizing the complexity of the issues, have refrained from rushing headlong into imposing new rules. Instead, governments have focused their efforts on a careful examination of the issues while also raising awareness of possible risks. For example, in a report entitled *Big Data: A Report on Algorithmic Systems, Opportunity, and Civil Rights*, the Obama administration examined how governments can best support growth in beneficial uses of big data while ensuring that it does not create or perpetuate discrimination.[42] While noting areas of concern in fields such as financial services, employment, higher education, and criminal justice, the report also recognized the potential to harness the power of big data to detect and eradicate such biases.[43] Several other policymakers have reached similar conclusions and recommendations, including the U.S. Federal Trade Commission (FTC).[44] Similarly, a growing concern among policymakers is the impact that widespread adoption of AI within enterprises could have on existing jobs, and specifically the risk that many workers in lower-skilled jobs could find themselves unemployed.

[42] See *White House Big Data Report, supra* note 14.

[43] See *id.* at 5.

[44] See *FTC Big Data Report, supra* note 2.

For instance, the *Future of AI* report notes that to the extent AI automates certain routine tasks, it will reduce demand for skills that can be automated while increasing demand for other skills that are complementary to AI.[45] But rather than regulating the use of technology, the report concluded that public policy can address these challenges by helping to ensure that workers are retrained and are able to succeed in occupations that complement, rather than compete with, automation.[46]

Even as governments consider these policy challenges, many are also forging ahead with plans to promote investment, innovation, and adoption of BDA and AI technologies. For instance, the Government of Canada has provided $125 million to fund a Pan-Canadian Artificial Intelligence Strategy focused on promoting research and developing talent, and intended to position Canada as a world leader in AI.[47] The Korean government has proposed an Artificial Intelligence Information Industry Development Strategy, which aims to strengthen the foundation for AI growth in Korea and released the report *Mid- to Long-Term Master Plan in Preparation for the Intelligent Information Society: Managing the Fourth Industrial Revolution*,[48] which provides in-depth analysis of artificial intelligence as a new driver of the economy.[49] Similarly, the Japanese government has announced a New Robot Strategy, designed to strengthen collaboration among industry, the government, and academia on matters related to robotics. It also issued a report entitled *The Conference on Networking Among AIs Report (2016): Impacts and Risks of AI Networking Issues for the Realization of Wisdom Network Society* (WINS), which offers the first systematic review of AI networking issues in Japan.[50]

For its part, Workday supports a differentiated and contextual approach to regulation in this area. Although Workday believes that certain general principles should apply across-the-board (for instance, that AI and related technologies should be deployed in ways that respect human dignity and rights), regulatory requirements should focus on specific applications and

[45] See *Preparing for the Future of AI Report,* available at https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf *supra* note 52, at 2.

[46] See *id.* at 29.

[47] *Canada funds $125 million Pan-Canadian Artificial Intelligence Strategy*, Cision (Mar. 22, 2017), available at http://www.newswire.ca/news-releases/canada-funds-125-million-pan-canadian-artificial-intelligence-strategy-616876434.html.

[48] Gov't of Korea, *Mid- to Long-Term Master Plan in Preparation for the Intelligent Information Society: Managing the Fourth Industrial Revolution*, available at http://www.msip.go.kr/dynamic/file/afieldfile/msse56/1352869/2017/07/20/Master%20Plan%20for%20the%20intelligent%20information%20society.pdf.

[49] Hyea Won Lee and Young-Jin Choi, *What Is Korea's Strategy to Manage the Implications of Artificial Intelligence?,* World Bank (Aug. 29, 2016), http://blogs.worldbank.org/ic4d/taxonomy/term/12749.

[50] See Fumio Shimpo, *Japan's Role in Establishing Standards for Artificial Intelligence Development*, Carnegie Endowment for International Peace (Jan. 12, 2017), http://carnegieendowment.org/2017/01/12/japan-s-role-in-establishing-standards-for- artificial-intelligence-development-pub-68311.

use scenarios and should be adopted only where there is concrete evidence of actual harm. Determining the appropriate contours of regulation likewise should take account of differences in the technologies involved (e.g., BDA vs. AI) and the contexts in which they are deployed (e.g., in the consumer vs. enterprise context).

For example, although both BDA and AI require thoughtful consideration about the data inputs used to train the relevant algorithms and that form the basis for specific decisions, to the extent that AI applications supplant decisions traditionally made by humans, these applications may raise additional questions about fairness and accountability, worker displacement, and unintended biases. Even where such concerns also apply to BDA, they may have less relevance in the enterprise context than in public-sector or consumer contexts—for instance, tools that help enterprises improve their supply-chain management are unlikely to implicate the same concerns that exist where such an application directly affects or interacts with consumers.

Throughout this debate, it is essential not to lose sight of the fact that all of us ultimately control how these technologies are developed and deployed. The fact that certain AI applications engage in limited types of automated decision-making and action does not detract from the fact that we as a community—of developers, users, consumers, and policymakers—have the power to ensure that we use them in accordance with our principles and values. This imposes important responsibilities on all of us—especially those of us in the technology sector—but also creates opportunities for the private sector to promote thoughtful, effective solutions to the policy issues highlighted above.

Fortunately, the private sector is rising to the challenge. For instance, an initiative on Fairness, Accountability, and Transparency in Machine Learning ("FAT-ML"), comprising academics and industry stakeholders, has articulated *Principles for Accountable Algorithms and a Social Impact Statement for Algorithms*.[51] These principles aim to help developers design, and users implement, algorithmic systems in publicly accountable ways that promote responsibility, explainability, accuracy, auditability, and fairness in the use of algorithms.[52]

[51] See *Principles for Accountable Algorithms and a Social Impact Statement for Algorithms*, http://www.fatml.org/resources/principles-for-accountable-algorithms (last visited July 27, 2016).

[52] See *id.*

In a similar vein, the nonprofit Center for Democracy and Technology has articulated an accountability framework for AI solutions based on the principles of fairness, explainability, auditability, and reliability. The framework is designed to "provide an alternative understanding of data-driven decisions, not as objective measures, but as tools that must be carefully calibrated in order to avoid biased outcomes."[53]

More recently, the Partnership on AI, a technology industry consortium focused on establishing best practices for AI systems, emphasizes the importance of developing best practices to create fair, explainable, and accountable AI systems. One of the Partnership's key themes is that when AI tools are used to supplement or replace human decision-making, those who develop and deploy them must ensure that they are safe, trustworthy, and aligned with the ethics and preferences of the people whose lives are affected by their actions.[54]

These and many other private-sector efforts—by industry, academics, civil society, and others—demonstrate both an awareness of the importance of the policy issues at stake, and a willingness to actively craft solutions. It is vital that policymakers take note of these efforts and give them time to develop and mature. In the section that follows, we consider the appropriate roles for both the public and private sector, and propose areas in which these groups can work together to find common solutions.

IV. The Path Forward

A. For the Private Sector

In thinking about the path forward, Workday believes that the responsibility to ensure that safe, ethical, and appropriate development and deployment of AI and related technologies lies first and foremost with industry. We believe the private sector has a particularly vital role to play in crafting principles, best practices, guidelines, and technology solutions in several specific areas. These include:

- Developing technologies, policies, and practices that effectively protect personal data and robustly secure data, systems, and devices

[53] Center for Democracy and Technology, *Digital Decisions,* available at https://cdt.org/issue/privacy-data/digital-decisions/ (last visited Oct. 9, 2017).

[54] *Thematic Pillars,* Partnership on AI to benefit people and society, available at https://www.partnershiponai.org/the matic- pillars/ (last visited Oct. 9, 2017).

- Committing to the responsible and ethical design of AI, including its fair and accountable use, and developing principles and guidelines to help generate a common understanding of what this means in practice
- Committing to respect human dignity and fundamental human rights in the deployment of AI and related technologies
- Undertaking further efforts to detect and eliminate unintended biases in AI and BDA technologies, including tools to help those who deploy these solutions avoid discriminatory outcomes
- Actively seeking inclusive solutions to help ensure that the benefits of BDA and AI are broadly shared across communities and geographies, and in particular benefit historically underserved populations

As a first step toward these goals, Workday believes the time is right for industry to develop a high-level code of conduct on BDA and AI development and deployment. In developing this code, industry should welcome and actively seek input from across all sectors of the economy and society. Given that AI and BDA increasingly will impact many aspects of people's lives across all walks of life, the code must likewise reflect and take into account the diverse interests of people and society.

To be effective, this code should articulate principles that will promote ethical and responsible action in the design of AI services and systems and throughout the lifecycle of AI implementations. These principles should include respect for fundamental human rights, detecting and eliminating unintended biases and discriminatory outcomes, protection of individual rights and freedoms, maintaining human control of AI, avoiding reasonably predictable misuses of AI, and ensuring human safety, security, and privacy.[55] This code should also recognize the paramount obligation of industry to comply with all laws and regulations that currently apply to these technologies.

Even as the code articulates cross-cutting principles, however, it should take into account that different technologies and deployment scenarios necessarily raise distinct concerns and risks. For instance, while the overarching principles that apply to AI and BDA will often be the same,

[55] See also World Economic Forum, *Top 9 Ethical Issues in Artificial Intelligence* (Oct. 21, 2016), https://www.weforum.org/agenda/2016/10/top-10-ethical-issues-in-artificial-intelligence/ (last visited Oct. 9, 2017).

the steps that developers and users must take to honor these principles might be more relevant to AI than BDA, given that BDA solutions only inform, and do not augment or supplant, human decision-making. Similarly, certain principles in the code might have far more relevance to AI solutions that directly interact with or impact individual consumers than to those used either solely within an enterprise (e.g., CRM analytics tools) or that are exclusively B2B (e.g., cross-enterprise inventory management tools).

Development of such a code cannot occur in isolation. It should be informed by research and input from technologists, scientists, ethicists, civil society, and others. It will also need to be accompanied by intensive dialogue, which will undoubtedly reveal differences in perspectives and goals, but will help ensure that the final code represents broad views and the greatest consensus possible. In addition, of course, the code will need to align with whatever developments occur in the applicable regulatory regimes and in policymaker discussions in this area.

B. For the Public Sector

Although the private-sector actions already discussed are critical to the responsible development and deployment of AI and related technologies, there are certain areas where governments can and should lead. These include:

- **EDUCATION AND SKILLS.** The growth of AI and BDA will dramatically increase the need for people with specialized skills. This includes not only AI researchers and engineers, but also a large number of specialists who refine AI methods for specific applications, and a much larger number of users who operate those applications in specific settings.[56] For researchers, AI training often requires a strong background in computer science, statistics, mathematical logic, and information theory, and for specialists, training typically requires a background in software engineering and in the application area.[57] Nations whose educational systems provide these skills are more likely to establish leading positions in key areas such as algorithm creation and development, capability demonstration, and commercialization.[58] Accordingly, a top priority for governments is to ensure that students are learning the skills they need to thrive in a world suffused with AI and BDA technologies. These efforts must begin in primary schools and extend throughout the educational system.

[56] See *Preparing for the Future of AI* Report, *supra* note 52, at 26.

[57] See *id.*

[58] See *NSTC AI Strategic Plan, supra* note 21, at 35.

- **WORKER RETRAINING**. Although many AI applications today merely replace tasks, not workers, and will increasingly create new jobs for people with highly valuable skills, it is also inevitable that some workers will be displaced. Because this displacement could happen relatively quickly, governments should prepare now to provide displaced workers with opportunities to develop new skills that will increase their chances of finding new jobs. Although further research is needed to better understand future workforce needs for AI and how best to fill such needs, governments that begin planning now for worker retraining will stand the best chance of avoiding the disruption and social unease that significant job displacement may generate.[59] Governments should also work to ensure that the benefits of AI are broadly shared and that it does not exacerbate existing inequalities.[60]
- **PROMOTING PUBLIC- AND PRIVATE-SECTOR RESEARCH.** Future breakthroughs in AI will require substantial investments in fundamental scientific research. Today, only a handful of companies have the resources and business incentives to invest in such research over the long time spans necessary to yield results. As in the past, governments can provide the "fuel" that ignites valuable practical applications of AI by investing in basic scientific research, then making the results of that research broadly available for commercialization. Governments should also adopt policies that provide incentives for private-sector R&D, including through innovative tax policies and by expanding access to financing, especially for start-ups and entrepreneurs.
- **ENSURING REGULATION IS FLEXIBLE, TECHNOLOGY-NEUTRAL, RISK-BASED, PROPORTIONATE, AND GROUNDED ON EVIDENCE.** Even as governments support private-sector solutions to the policy challenges raised by AI, it is inevitable that they will at times need to regulate. In doing so, they should ensure that such regulation is flexible and technology-neutral, so that companies are not forced into technical or other mandates that could freeze innovation and instead offer users multiple competing solutions that achieve the regulation's goals. So that regulation does not paint with too broad a brush, regulators should also take care to ensure that rules are risk-based, proportionate to the concrete harms at issue, and grounded on solid evidence.

[59] See *id.* at 35–36.

[60] See *Preparing for the Future of AI* Report, *supra* note 52, at 30.

V. Conclusion

Today we are poised on the cusp of the next technological revolution, driven by BDA and AI. Although this revolution is likely to make itself felt most immediately and profoundly in the workplace, it will increasingly permeate many different aspects of life and all sectors of the economy and society. The potential benefits are enormous; addressing the potential challenges that may arise without undermining these benefits requires thoughtful and deliberate policymaking.

As a first step in this journey, Workday hopes that this paper helps to ignite a dialogue. This dialogue should include not only policymakers and industry, but also data scientists, academics, civil society, and all others who may influence, or whose lives may be influenced by, the development and deployment of these technologies. Only a broad conversation will ensure that all voices are heard and that all views are considered. We hope that a national—and ideally global—conversation will initiate the development of a code of conduct to help guide action in this area. To be truly useful, this code must acknowledge the differences between the technologies themselves as well as the differences in the contexts in which they are deployed. The secondary and long-term effects of big data and AI technology can be far-reaching indeed, but they are not forces beyond our control. It is incumbent on us to properly manage their growth and unlock their tremendous potential to improve all aspects of society.

Workday | Phone: +1 (925) 951-9000 | workday.com